Brief Announcement: Proximal Byzantine Agreement: Improved Accuracy for Fault-Tolerant Replicated Datastreams

Roy Shadmon
University of California, Santa Cruz, CA, USA

Owen Arden
University of California, Santa Cruz, CA, USA

Abstract

Approximate Byzantine Agreement (ABA) protocols enable nonfaulty replicas with different initial values to derive a values within a ϵ-neighborhood of each other, despite the presence of Byzantine behavior. While they give strong guarantees for this ϵ-*agreement* property, they tend to have weaker guarantees that the derived value is *accurate* with respect to some ground truth. Worse, they often have impractical requirements such as large replica sets proportional to data dimensionality, or a priori knowledge of the maximum distance between nonfaulty values.

In *Stochastic Byzantine Agreement (SBA)*, the distribution of the nonfaulty values is the result of a stochastic process influenced by sensor measurement error or other sources of noise that affect system outputs. For these scenarios, we present *Proximal Byzantine Agreement* (PBA), a stochastic Byzantine agreement protocol which infers the most likely output of replicated computation based on the proposed values observed by each replica. Unlike ABA protocols, PBA prioritizes accuracy over agreement. PBA accuracy is relative to the variance of nonfaulty values, yielding comparatively more accurate results for noisy data, particularly when noise is asymmetric. Our evaluations demonstrate this accuracy scales with data dimensionality, outperforming or only mildly underperforming methods that require quorums with up to 10× *more replicas* and 4× *to* 124× *more computation time* per agreement decision, even at relatively low dimensions ($d = 4$ to $d = 18$).

2012 ACM Subject Classification Computer systems organization → Dependable and fault-tolerant systems and networks; Computer systems organization → Sensor networks

Keywords and phrases Byzantine fault tolerance, distributed control systems, robust statistics

Digital Object Identifier 10.4230/LIPIcs.DISC.2025.64

Funding Partial funding provided by 2025 Berkeley RDI AI & Decentralization Innovation Award.

1 Introduction

Reliably controlling systems and devices that interact with the physical environment is challenging in any context, but especially in edge networks. Volatile communication, power, and operating conditions demand greater fault tolerance, and edge devices and sensors are often more vulnerable to malicious manipulation than those in secure, monitored facilities. These realities complicate the coherent design of sensor data-processing pipelines, distributed control systems, and multi-agent robotics applications – scenarios that are *functionally* similar but *operationally* require ad hoc mitigation techniques to tolerate deployment conditions.

The growing demand for such systems spans consumer-facing devices such as home-automation sensors and controls, personal health monitoring devices and drug-delivery systems, and automotive telemetry and safety systems, as well as industrial and agricultural platforms for monitoring and controlling mechanical, physical, and biological processes. In all of these domains, the quality and reliability of sensor and sensor-derived data is central

to effective operation. When that data is delayed, missing, or corrupted – whether from environmental noise, hardware faults, or malicious interference – the consequences can range from suboptimal performance to catastrophic, even unsafe, failures.

Despite shared characteristics, a divergence occurred between communities that prioritize agreement "regardless of what they agree upon" originally proposed by Lamport [9], and communities that also want guarantees on the accuracy of some quantity being approximated. Whereas most agreement protocols give *validity* [5] specifications on their output in terms of the initial values proposed by nonfaulty replicas, *accuracy* [12] specifications are with respect to a true quantity approximated by the outputs. Much subsequent work on approximate agreement does not explicitly consider accuracy, with some notable exceptions [1, 3] in the sensor fusion community. In all cases, relatively strong assumptions are made about nonfaulty values such as the maximum pairwise distance [10], maximum width of a containing interval [12], or the true value is contained in the intersection of ranges of nonfaulty replicas [3]. Violating these assumptions affects accuracy and can lead to arbitrary protocol failure.

This paper restores focus on the accuracy of an important subset of approximate Byzantine agreement (ABA) problems under weaker, more realistic assumptions. Since phenomena measured by sensors can often be modeled as stochastic processes [17], we define *Stochastic Byzantine Agreement (SBA)*, a refinement of the ABA problem. In the SBA problem, "true" values are generated by a stochastic process with an unknown *output distribution*, and the observations of those values proposed by nonfaulty replicas are perturbed by noise drawn from an unknown *error distribution*. Byzantine replicas act arbitrarily, and thus their proposed values are not generated – or even effectively approximated – by any distribution. Our SBA protocol, Proximal Byzantine Agreement (PBA), uses a novel *value selection process*, how replicas select outputs based on the messages they receive, that borrows techniques from robust statistics [7]. As long as a majority of these messages are stochastically distributed, the influence of arbitrarily selected values on the inference process is limited to a fixed bound.

This brief announcement reports on progress following our initial evaluation [18] of the feasibility of this value selection technique. Previously, only one-dimensional, normally distributed distributions were explored, and probabilistic protocol properties were described informally. Here we give formal, geometric definitions and theorems with rigorous guarantees, including evaluations comparing PBA to multiple ABA protocols on multidimensional data. These evaluations indicate a surprising degree of scalability with respect to dimensionality: only minor losses in accuracy compared to ABA protocols that require $10\times$ *more replicas* and $4\times$ *to* $124\times$ *more computation time* per agreement decision, even at relatively low dimensions ($d = 4$ to $d = 18$). Since each additional output offers more information from the unknown output distribution, our results indicate PBA can extract more information from observations, reinforcing the tradeoff [12] between high-precision agreement and statistical accuracy.

2 Proximal agreement as a probability maximization

Replicas participating in Approximate Byzantine Agreement (ABA) protocols (e.g.,[5, 14, 19]) start with an arbitrary real value and reach an output solution satisfying two properties:

1. ϵ-*Agreement*: all nonfaulty replicas eventually output values within ϵ of each other.
2. *Validity*: these values are bound (in some sense) by the initial values of nonfaulty replicas.

Mahaney and Schneider [12] distinguish Inexact Agreement (IA) from classical ABA by redefining the validity property in terms of *accuracy*. Outputs are instead bound by their distance to a true value $\hat{v}$, which each nonfaulty replica's initial value v_i approximates.

Our work concerns a new class of protocols that make an additional assumption, often satisfied by the underlying data processed by IA systems. Stochastic Byzantine Agreement (SBA) protocols require that $\hat{v}$ and nonfaulty v_i are – or can be effectively approximated as being – distributed according to probability distributions. No assumptions are made about the nature of Byzantine values. The two distributions are distinct and potentially unknown, but correlated since v_is approximate $\hat{v}$. SBA protocols also satisfy ϵ-agreement, but ϵ in SBA is larger than the ϵ values in classical ABA protocols. This difference in ϵ-agreement is by design and unavoidable due to tradeoffs between accuracy and high-precision agreement [12]. The goals of these protocols are close enough that we argue ABA is an appropriate umbrella term, with IA and SBA refining properties 1 and 2 as described above.

The key characteristic of our SBA protocol, *Proximal Byzantine Agreement* (PBA), is that each replica uses the set R of received values $v_i \in R$ to find a quorum of values r and a candidate $\hat{v}$ such that $\hat{v}$ would be the most likely true value if we knew that the quorum values were nonfaulty. Selecting the highest-probability $\hat{v}$ conditioned on the observations in r, ensures that if any (minority of) values in r turn out to be faulty, then the quorum r (and thus the inferred $\hat{v}$) is at least a likely as a quorum containing only nonfaulty values. At a high level, this process is described by the probability maximization in Def. 2.1.

▶ **Definition 2.1** (Proximal Byzantine Value Selection). *For received values $R \subseteq \mathbb{R}^n$ and quorum size s, select a set r from the s-sized subsets of R, and a value x that maximizes the likelihood that x is the true value given observations in r.*

$$\mathsf{PBA}(R, s) \triangleq \underset{\substack{r \in [R]^s \\ x \in \mathbb{R}^n}}{\operatorname{argmax}} P(x \mid r)$$

Our chosen approximation of Proximal Byzantine Value Selection (PBVS) uses the geometric median of each r as a robust estimator for the expected value of the unknown distribution of observations in r, thus approximating the candidate x. Below we will just use $\mathsf{PBA}(R, s)$ to refer to our approximation of PBVS to avoid additional notation. The geometric median [16] is well-suited for estimating the expected value of a broad class of distributions [15, 16], and remains robust even when $< \frac{1}{2}$ of the samples are arbitrarily corrupted. We compare the conditional likelihoods $P(x \mid r)$ for each r and x with a similarity-based approximation of conditional probabilities [2]. This is necessary for multidimensional data since the covariance matrix is unknown, but is a reasonable approach since each nonfaulty replica output is an approximation of the same underlying value.

More specifically, for each quorum r, we compute its median $x = \textsc{Gm}(r)$ and evaluate the likelihood $\widehat{P}(x \mid r)$ using a similarity measure adapted from Blok et al. [2]. $\widehat{P}(x \mid r)$ approximates the affect of the unknown covariance matrix and joint probability of q on the independent probability of x with an exponent α to score similarity between x and quorum values on a $[0, 1)$ scale. We chose this approach after discovering that computing the geometric median and then approximating its likelihood is significantly faster than iteratively evaluating $\widehat{P}(X = x \mid r)$ to find maximal values, and had almost no impact on accuracy.

2.1 Geometric properties

Approximation of a true value introduces a degree of uncertainty for systems that use those approximations to make decisions. We are unaware of any ABA protocols that estimates the uncertainty of its outputs.[1] In many scenarios, it may be better to perform no action at all than to act on very uncertain data. Therefore, calculating a concrete guarantee on each output is an important feature of PBA protocols.

[1] In a sense, the strong assumptions of existing IA protocols [12, 4, 3] act as *inputs* about uncertainty.

A well-known result [11] is that the displacement of the geometric median $\text{GM}(Q)$ of a set of nonfaulty values is bounded when up to $|F| < |Q|$ faulty values are added to the set.

▶ **Lemma 2.2** (Theoretical maximum displacement [6, 11]). *Let $R = Q \cup F$, $|Q| = s$ and $|F| = f$ such that $s > f$ with geometric medians $\text{GM}(Q) = m_Q$ and $\text{GM}(R) = m_R$. Let $\Delta^Q = \max_{q_i \in Q} ||q_i - \text{GM}(Q)||_2$ and $C_0^{(s,f)} = \frac{s}{\sqrt{s^2 - f^2}}$. Then the maximum displacement is*

$$||m_Q - m_R||_2 \leq C_0^{(s,f)} \Delta^Q$$

However, the contents of Q is unknown, so the bound implied by Lemma 2.2 cannot be directly calculated from the information available. Instead, Lemma 2.3 presents a computable bound guaranteed to contain the geometric median of the nonfaulty values, but is based only on the received values and the system parameters (i.e., n, f, and s).

▶ **Lemma 2.3** (Computable maximum displacement). *Let $R = Q \cup F$, $|Q| = s$ and $|F| = f$ such that $s > 2f$ with $\text{GM}(Q) = m_Q$ and $\text{GM}(R) = m_R$, and $C^{(s,f)} = \frac{s}{\sqrt{2sf - f^2}}$. Then*

$$||m_Q - m_R||_2 \leq C^{(s,f)} \Delta_f^R$$

The computable displacement is used to compute the *region guarantee* for a given PBA result. The fact that PBA only produces outputs with region guarantees containing the true output is our instantiation of an *accuracy* [12] property.

▶ **Definition 2.4** (Region Guarantee). *For any set $R \subseteq \mathbb{R}^d$, $|R| > s + f$ where $s > f$, define $R^{(f,x)} \subseteq R$ by removing the f elements of R furthest from x. Let $\Delta_f^R = \max_{r_i \in R^{(f, \text{GM}(R))}} ||r_i - \text{GM}(R)||_2$. The* region guarantee *$RG(R, s, f)$ is a d-dimensional ball centered at x with radius $2C^{(s,f)}\Delta_f^R$.*

▶ **Lemma 2.5** (Accuracy). *For any $R = (Q \cup F)$ containing $|Q| = s > 2f$ nonfaulty and at most $|F| = f$ faulty values, let $\mathsf{PBA}(R, s) = (x, R')$. Then $\text{GM}(Q) \in RG(R, s, f)$ and*

$$||m_Q - x||_2 \leq C_0^{(s,f)} \Delta^{R'} + C_0^{(s,f)} \Delta^Q \leq 2C^{(s,f)} \Delta_f^R$$

2.2 OneShot PBA Agreement and Termination

For this brief announcement, we describe OneShot PBA, a protocol that prioritizes accuracy. It infers the potential distances between the true value, its output, and other replica outputs using only the messages it receives, without additional coordination. Clients can receive proposed values directly from replicas and execute our PBVS algorithm locally to select values. Like most IA protocols [3, 4, 13], but unlike other ABA protocols [5, 12, 14, 19], OneShot does not iterate to converge on a single answer. Doing so would obscure the initial values, sacrifice accuracy, and undermine the fidelity of the region guarantees.

The termination conditions are thus straightforward: OneShot PBA requires $s + f$ values where $s \geq 2f + 1$, so termination for the asynchronous case requires at least $n \geq s + 2f \geq 4f + 1$ replicas. Otherwise, f faulty replicas could withhold values without being distinguished from f nonfaulty replicas experiencing a network partition, and halt progress.

PBA OneShot's requirements on replica set size compare favorably to other ABA protocols (Table 1), particularly when considering its lack of dependency on data dimensionality and weak assumptions regarding the range of correct values. Unlike prior work that tend to improve accuracy with more replicas, a distinguishing characteristic of PBA value selection is that more replicas are not required to ensure the safety or liveness properties of the system, yet still benefits from the increase in accuracy.

Table 1 Replica set size. d is number of data dimensions.

Algorithm	$n \geq$
OneShot	$4f + 1$
ABA [5]	$5f + 1$
IA [12]	$3f + 1$
BI [3, 4]	$2df + 1$
BVC [14, 19]	$(d + 2)f + 1$

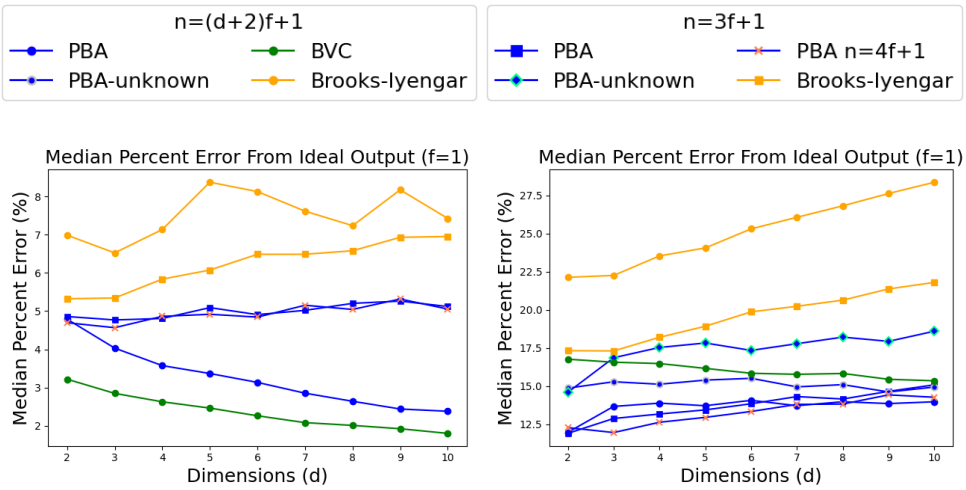

(a) Gaussian samples, Byzantine attack. **(b)** Exponential samples, Byzantine attack.

Figure 1 Empirical simulations measuring PBA's accuracy compared to baseline algorithms given data distributed by a Gaussian (symmetric) and Exponential (asymmetric) distributions.

Despite the absence of inter-replica coordination or iteration, the PBA values derived by consumers from different subsets are bounded with respect to their distance to the true expected value. This bound is in terms of the displacement properties of the geometric median discussed in §2.1, and contingent on quorum sizes representing a sufficiently large fraction of the nonfaulty replicas, in this case $> \frac{|Q|}{2}$.

It is taken for granted in other protocols that quorums must contain at least half of all nonfaulty replicas since these protocols ensure all quorums intersect. We only make this size requirement explicit for agreement because not all instances of PBA necessarily involve multiple consumers. For example, there is no need for agreement if a system controller is the only consumer of replica outputs, as is often the case in industrial control systems. For contrast, note that accuracy (Lemma 2.5) only requires quorum membership to be more than twice the number of faulty replicas.

▶ **Lemma 2.6** (Agreement). *For any* $R^i = (Q^i \cup F^i) \subseteq R = (Q \cup F)$, *let* $r = |R^i|$, $s = |Q^i|$, $f = |F|$, *and* $\mathit{PBA}(R^i, s) = (x_i, R^i_*)$. *If* $s > \left\lceil \frac{|R|}{2} \right\rceil$ *and* $r \geq s + f$, *then for any other subset* $R^j = (Q^j \cup F^j) \subseteq R$ *where* $|R^j| \geq r$ *and* $\mathit{PBA}(R^j, s) = (x_j, R^j_*)$, $||x_i - x_j|| \leq 2C_0^{(s,f)}\Delta^Q$.

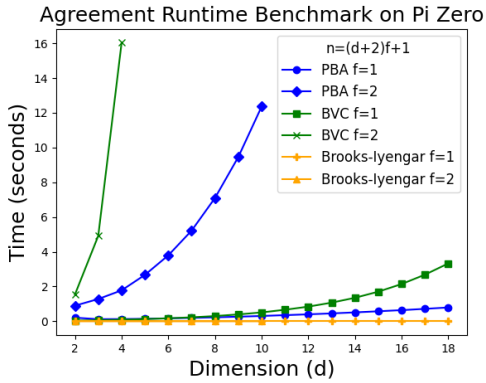

Figure 2 Runtime benchmark (Pi Zero).

3 Evaluation

PBA accuracy compared to baseline. We perform repeated simulations on Chameleon Cloud [8], comparing PBA under several quorum sizes against approximate Byzantine Vector Consensus (BVC) [19] and Brooks-Iyengar (BI) [4]. Here we present only the Byzantine attack scenario. Each trial samples $n - f$ d-dimensional replica outputs x_i from a random Gaussian or Exponential distribution (with a fixed, low variance) and multiplies each element-wise by a noise vector y_i independently sampled from an error distribution with $E[Y] = 1$ and (uniformly-sampled) standard deviation from 0.01 to 0.09. We compute the agreement vector of each protocol and report the median percent error with respect to x_i across 1000 trials.

Figure 1a compares results when $n - f$ outputs are sampled from a Gaussian distribution under an optimal Byzantine attack with $f = 1$. While BVC remains highly robust, PBA achieves comparable accuracy with just $n = 4f + 1 = 5$ (small n) outputs, versus BVC's $n = (d+2)f + 1 = 13$ (large n) outputs. This indicates that at least some of BVC's accuracy may result from its reliance on more samples, but any ABA protocol will improve in accuracy with more samples: aggregating values concentrates outputs around the true value. PBA approaches the true value with fewer samples (i.e., replicas) and performance overhead.

In the Exponential and large n setting ($n = (d+2)f+1$), Figure 1b shows PBA outperforms all other $n = 4$ protocols. In the small n setting, PBA demonstrates superior accuracy to all baselines – even in the PBA-unknown case, where the prior distribution is estimated as Gaussian. This accuracy advantage likely stems from PBA's quorum selection strategy and agreement inference method. In contrast, the benefit other protocols received from additional samples in the Gaussian case, a symmetric distribution, are tempered in the exponential case, an asymmetric distribution. The gap in accuracy suggests the averaging effect occurring in other protocols generalizes poorly to other distributions, whereas PBA's more statistically-informed approach adapts more easily.

Runtime performance. Figure 2 presents runtime benchmarks for $f = 1$ and $f = 2$ on a node hosted on a single Raspberry Pi Zero 2W. These results are based on repeated simulations of the PBA, BVC, and BI protocols operating on sets of n d-dimensional vectors. where $n = (d+2)f + 1$, and PBA quorum sizes set to $s = n - f$. PBA is faster than BVC for $f = 1$, $d > 5$ (with similar performance at lower d), and more than $4\times$ faster at $f = 1$, $d = 18$. In all $f = 2$, $d = 4$ scenarios, PBA significantly outperforms BVC ($124\times$). While BI also exhibits scalable runtime in high-dimensional settings, each replica must produce outputs

with a known error range that contains the true value, and that must also intersect the ranges of all other nonfaulty nodes. If these invariants are violated by any nonfaulty replica, the system may not withstand Byzantine attack. Even without an attack, the protocol could fail to terminate.

Our empirical results against BVC reinforces that PBA offers a compelling trade-off: it maintains competitive accuracy while substantially reducing computational cost. Beyond performance, PBA also provides robust fault tolerance and accurate agreement with significantly fewer replica outputs. Notably, PBA's minimum required number of replicas and quorum size are independent of d; in contrast to BVC, whose replica requirements scale significantly with d. Under-provisioned PBA systems instead have more uncertainty, but the amount is characterized by region guarantees. This scalable perfomance and graceful, principled degradation indicate promise for both large-scale and resource-constrained sensor-based control systems and datastream processing applications.

References

1 Buke Ao, Yongcai Wang, Lu Yu, Richard R. Brooks, and S. S. Iyengar. On Precision Bound of Distributed Fault-Tolerant Sensor Fusion Algorithms. *ACM Computing Surveys (CSUR)*, May 2016. doi:10.1145/2898984.

2 Sergey Blok, Douglas Medin, and Daniel Osherson. Probability from similarity. In *AAAI Spring Symposium on Logical Formalization of Commonsense Reasoning*, pages 36–42, 2003.

3 Richard Ree Brooks and S. Sitharama Iyengar. Optimal matching algorithm for multidimensional sensor readings. In *Sensor Fusion and Networked Robotics VIII*, volume 2589. SPIE, September 1995. doi:10.1117/12.220948.

4 R.R. Brooks and S.S. Iyengar. Robust distributed computing and sensing algorithm. *Computer*, 29(6), June 1996. doi:10.1109/2.507632.

5 Danny Dolev, Nancy A. Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. Reaching approximate agreement in the presence of faults. *Journal of The Acm*, 33(3), May 1986. doi:10.1145/5925.5931.

6 El-Mahdi El-Mhamdi, Sadegh Farhadkhani, Rachid Guerraoui, and Lê-Nguyên Hoang. On the Strategyproofness of the Geometric Median. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*. PMLR, April 2023. URL: https://proceedings.mlr.press/v206/el-mhamdi23a.html.

7 Peter J. Huber and Elvezio M. Ronchetti. *Robust Statistics*. Wiley, 2011.

8 Kate Keahey, Jason Anderson, Zhuo Zhen, Pierre Riteau, Paul Ruth, Dan Stanzione, Mert Cevik, Jacob Colleran, Haryadi S. Gunawi, Cody Hammock, Joe Mambretti, Alexander Barnes, François Halbach, Alex Rocha, and Joe Stubbs. Lessons learned from the chameleon testbed. In *Proceedings of the 2020 USENIX Annual Technical Conference (USENIX ATC '20)*. USENIX Association, July 2020. URL: https://www.usenix.org/conference/atc20/presentation/keahey.

9 Leslie Lamport. The weak Byzantine generals problem. *Journal of the ACM*, 30(3), 1983. doi:10.1145/2402.322398.

10 Leslie Lamport and P. M. Melliar-Smith. Byzantine clock synchronization. In *Principles of Distributed Computing (PODC)*, 1984. doi:10.1145/800222.806737.

11 Hendrik P. Lopuhaa and Peter J. Rousseeuw. Breakdown Points of Affine Equivariant Estimators of Multivariate Location and Covariance Matrices. *The Annals of Statistics*, 19(1), March 1991. doi:10.1214/aos/1176347978.

12 Stephen R. Mahaney and Fred B. Schneider. Inexact agreement: Accuracy, precision, and graceful degradation. In *Principles of Distributed Computing (PODC)*. ACM Press, 1985. doi:10.1145/323596.323618.

13 Keith Marzullo. Tolerating failures of continuous-valued sensors. *ACM Transactions on Computer Systems*, 8(4), November 1990. doi:10.1145/128733.128735.

14 Hammurabi Mendes and Maurice Herlihy. Multidimensional approximate agreement in Byzantine asynchronous systems. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*. ACM, June 2013. doi:10.1145/2488608.2488657.

15 Stanislav Minsker. Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21(4), November 2015. doi:10.3150/14-BEJ645.

16 Stanislav Minsker and Nate Strawn. The Geometric Median and Applications to Robust Mean Estimation. *SIAM Journal on Mathematics of Data Science*, 6(2), June 2024. doi:10.1137/23M1592420.

17 Emanuel Parzen. *Stochastic Processes*. Dover Books on Mathematics. Dover Publications Inc., 2015.

18 Roy Shadmon and Owen Arden. Enhancing accuracy in approximate byzantine agreement with bayesian inference. In *55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks - Supplemental Volume (DSN-S)*, 2025. doi:10.1109/DSN-S65789.2025.00064.

19 Nitin H Vaidya and Vijay K Garg. Byzantine vector consensus in complete graphs. In *Proc. of the 2013 ACM Symposium on Principles of Distributed Computing*, 2013.